

February 26, 2019

RT0982

Artificial Intelligence 11 pages

Research Report

Real-time tree search with pessimistic scenarios

Takayuki Osogami
Toshihiro Takahashi

IBM Research – Tokyo
IBM Japan, Ltd.
19-21 Nihonbashi Hakozaki-cho, Chuo-ku
Tokyo 103-8510, Japan

Real-time tree search with pessimistic scenarios

Takayuki Osogami

Toshihiro Takahashi

Abstract

Autonomous agents need to make decisions in a sequential manner, under partially observable environment, and in consideration of how other agents behave. In critical situations, such decisions need to be made in real time for example to avoid collisions and recover to safe conditions. We propose a technique of tree search where a deterministic and pessimistic scenario is used after a specified depth. Because there is no branching with the deterministic scenario, the proposed technique allows us to take into account far ahead in the future in real time. The effectiveness of the proposed technique is demonstrated in Pommerman, a multi-agent environment used in a NeurIPS 2018 competition, where the agents that implement the proposed technique have won the first and third places.

1 Introduction

Autonomous agents, such as self-driving cars and drones, need to make decisions in real time, which is particularly important but difficult in critical situations for example to avoid collisions. Such decisions often need to be made in a sequential manner to achieve the eventual goal (*e.g.*, avoiding collisions and recovering to safe conditions), under partially observable environment, and by taking into account how other agents behave. Towards this far-reaching goal of realizing such autonomous agents, we propose practical techniques of sequential decision making in real time and demonstrate their effectiveness in Pommerman, a multi-agent environment that has been used in one of the competitions held at the Thirty-second Conference on Neural Information Processing Systems (NeurIPS 2018) on Dec. 8, 2018 Resnick *et al.* [2018a]. The techniques that we propose in this paper have been used in the Pommerman agents (HakozakiJunctions and dypm-final) who have won the first and third places in the competition.

In Pommerman, a team of two agents competes against another team of two agents on a board of 11×11 grids (see Figure 1 (a) for an initial configuration of the board). Each agent can observe only a limited area of the board, and the agents cannot communicate with each other. The goal of a team is to knock down all of the opponents. Towards this goal, the agents place bombs to destroy wooden walls and collect power-up items that might appear from those wooden walls, while avoiding flames and attacking opponents. See Figure 1 (b) for an example of the board in the middle of the game. See Resnick *et al.* [2018a] and the GitHub repository¹ for details of Pommerman.

Although Pommerman has been developed recently (initial GitHub commit was Dec. 25, 2017), it has been gaining much attention as a benchmark of multi-agent study in the field of planning, game theory, and reinforcement learning Kapoor [2018]; Hernandez-Leal *et al.* [2018]. Prior to the one at NeurIPS 2018, the first competition was held on Jun. 3, 2018. The winning agent sets an intermediate goal with heuristics and performs depth-limited tree search to achieve that intermediate goal Zhou *et al.* [2018]. Resnick *et al.* [2018b] propose a technique of imitation learning with curriculum and show that, by using the behavioral data of the winning agent, it can train an agent as strong as the winning agent.

While the community has made progress in developing strong agents for Pommerman, the level of those agents is not yet comparable to what has been achieved for backgammon Tesauro [1994], Chess Campbell *et al.*, Atari video games Mnih *et al.* [2015], Poker Bowling *et al.* [2015]; Brown and

¹<https://github.com/MultiAgentLearning/playground>

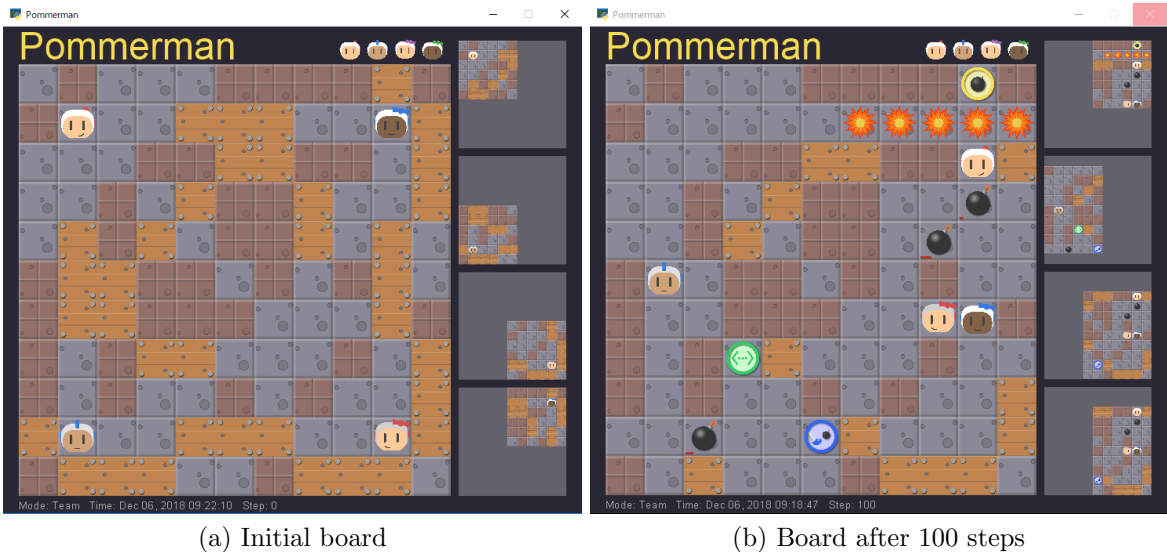


Figure 1: An initial board (a) and a board after 100 steps in Pommerman. The four small windows on the right most column respectively denote the areas that the four agents can observe.

Sandholm [2017], and Go Silver *et al.* [2017]. Pommerman has its own difficulty that prohibits effective applications of existing approaches that have seen success in other games.

What makes Pommerman difficult is the constraint on real-time decision making (*i.e.*, an agent needs to choose an action in 100 milliseconds). This constraint significantly limits the applicability of Monte Carlo Tree Search, which would otherwise be a reasonable approach to Pommerman Matiesen [2018]. In Pommerman, the branching factor at each step can be as large as $6^4 = 1,296$, because four agents take actions simultaneously in each step, and there are six possible actions for each agent. The agents should plan ahead and choose actions by taking into account the explosion of bombs, whose lifetime is 10 steps. Tree search with insufficient depth (less than 10) would ignore the explosion of bombs, which in turn would make the agents easily caught up in flames. Tree search with sufficient depth (at least 10) is practically infeasible with the large branching factor. Other difficulties of Pommerman include the following. Reward is only given at the end of an episode, which can be as long as 800 steps. The agents can observe only a limited part of the board, and some of the key information cannot be directly observed. The agent needs to coordinate with its teammate without explicit communication.

Here, we propose a practical approach to real-time tree search that allows us to take into account critical events far ahead in the future. In our approach, tree search after a specified depth is performed under the assumption of a deterministic and pessimistic scenario (*i.e.*, sequence of states). Because the scenario is deterministic, there is no branching after the specified depth, which allows us to perform the tree search with sufficient depth to take into account the critical events far ahead in the future. This deterministic scenario is designed to be pessimistic by allowing multiple unfavorable events can happen (*e.g.*, by letting opponents take multiple actions) simultaneously in a nondeterministic manner. Hence, our pessimistic scenarios are unrealistic in general. Our key idea is that an unrealistic scenario can capture critical events in the future better than a small number of realistic scenarios that can be sampled and explored under the constraint of real-time tree search. We adjust the level of pessimism via self-play to achieve the best overall performance.

Our approach is proposed particularly for Pommerman but might be generally applicable to other domains that require sequential decision making in real time. We demonstrate the flexibility of the proposed approach by instantiating it as two variants of Pommerman agents, who need to deal with the complex environment that involves multiple agents and partial observability. The effectiveness of the proposed approach is shown with Pommerman. The new approach of real-time tree search with

deterministic and pessimistic scenarios and its application to Pommerman constitute the contributions of this paper.

2 Related Work

There has been a significant amount of work on the techniques for real-time tree search particularly for real-time (strategy) games. As we will discuss it in the following, however, the focus of the prior work is on the techniques for reducing the search space or guiding the search towards the most relevant subspace. The novelty in our approach is in synthesizing the deterministic and pessimistic scenarios.

The prior work has investigated various techniques to make Monte Carlo Tree Search (MCTS) applicable to real-time games such as Ms. Pac-Man Pepels *et al.* [2014]; Ikehata and Ito [2011], StarCraft Uriarte and Ontañón [2016], Wargus Balla and Fern [2009], Physical Traveling Salesman Problem Powley *et al.* [2012], Quantified Constraint Satisfaction Problem Baba *et al.* [2011], and μ RTS Barriga *et al.* [2018]; Mariño *et al.* [2019]. An example of a recent work in this line is Mariño *et al.* [2019], who study a technique of action abstraction and apply it to MCTS among others to reduce the search space. Barriga *et al.* [2018] study a technique of using non-deterministic rules to reduce the branching in MCTS. In all of such prior work, MCTS is performed only with realistic or legal moves.

Guo *et al.* [2014] study the approach of using MCTS to generate training data for learning a deep neural network that approximates a policy or a value function, and the effectiveness of the proposed approach is demonstrated in Atari games. This approach is motivated by the observation that tree search (planning-based approaches) can perform far better than model-free approaches if it were not for the real-time constraint.

Real-time tree search has also been studied for deterministic settings. Here, the search tree is expanded on the basis of heuristic values as long as time permits. Similar to real-time MCTS, key questions are where to expand and what actions to take given the search tree investigated. For example, Mitchell *et al.* [2019] propose a risk-sensitive approach to these questions.

Our approach of synthesizing a pessimistic scenario is also related to the null-move heuristic, which has been studied particularly for Chess Goetsch and Campbell [1990], in that it considers a scenario with illegal moves. The null-move heuristic assumes that a player skips a move, which is illegal in chess, to estimate a lower bound of the value of the best move. The lower bound is then used to prune the search space. In contrast, our approach can assume that a player takes multiple moves in the pessimistic scenario.

3 Real-time tree search for Pommerman

In Pommerman, the dynamics of the environment is known, and much of the uncertainties resulting from partial observability can be resolved with careful analysis of historical observations. MCTS would thus be a competitive approach if it were not for the real-time constraint. For example, consider a situation where an agent can survive only by following a particular route. Tree search is particularly suitable for finding such a route, while model-free approaches of learning policies or value functions, if not impossible, would require large scale functional approximators (*e.g.*, deep neural networks) and a large amount of data for training to be able to follow that route. The applicability of MCTS or tree search in general is however significantly limited in Pommerman due to the real-time constraint and the large branching factor.

One approach of tree search is to push the depth as far as possible, and this is the approach taken by the `gorogm_eisenach` (`eisenach`) agents, who won the second place in the NeurIPS 2018 Pommerman competition. The `eisenach` agent was implemented in C++ with various engineering tricks to achieve the average depth of 2 in the tree search Resnick *et al.* [2019].




Figure 2: Tree search with deterministic and pessimistic scenarios.

3.1 Tree search with pessimistic scenarios

In our approach, the tree search after a specified depth is performed under a deterministic and pessimistic scenario, which will be simply referred to as a deterministic scenario or a pessimistic scenario in the following. Figure 2 shows an example. Here, the tree search is performed in a standard manner until the depth of 2. In this example, the branching factor is 2, and there are 4 nodes at the depth of 2. From each of these 4 nodes, “tree search” is continued until it reaches the depth of 5 by assuming a deterministic scenario. Because the scenario is deterministic, there are no branches after the depth of 2.

One may also interpret our approach as evaluating the 4 leaves (at the depth of 2) on the basis of the deterministic scenario from the depth of 2 to the depth of 5, and our following discussion will be based on this interpretation. Our approach keeps the size of the search tree small, because there are branches only until a limited depth. At the same time, our approach can take into account far ahead in the future, because the leaves (at the limited depth) can be evaluated with a deterministic scenario that can be much longer than what would be possible with branches. What differs from the rollout in MCTS is that we let the deterministic scenario be pessimistic, as we will discuss in Section 3.2.

More specifically, in the Pommerman agents that implement the proposed approach (*i.e.*, `Hakozaki-Junctions` (`hakozaki`) and `dypm-final` (`dypm`)), the depth of tree search is limited to one. The `hakozaki` agent considers all of the leaves at depth 1 but might need to choose an action before exhaustively searching all of the leaves due to timeout. On the other hand, the `dypm` agent considers six leaves at depth 1 by taking into account only its own actions, and the effect of the actions by the other agents are taken into account in evaluation with the deterministic scenario. In both of `hakozaki` and `dypm`, the leaves are evaluated with a deterministic scenario with the length of at least 10 to take into account the explosion of bombs, whose lifetime is 10. Recall that the `eisenach` agent can perform the tree search only at the average depth of 2. Hence, if we performed the standard tree search for 2 steps, there would be no computational budget left for the evaluation with deterministic scenarios.

3.2 Generating pessimistic scenarios

From each of the leaves in the search tree, we generate a deterministic scenario. The deterministic scenario is then used to estimate the value of each leaf, which in turn is used to determine the next best action of the agent. A key idea in our approach is to make this deterministic scenario be pessimistic, which we will discuss in this section with specific instantiation as Pommerman agents.

A pessimistic scenario can be generated in a systematic manner as follows. We assume that the state of the environment can be represented by the positions of objects. In Pommerman, these objects are agents, bombs, flames, powerup items, and walls. Some of those objects change their positions

randomly or by depending on the actions of the agents, which forces tree search to have branches. If one can tell the worst sequence of the positions of an object among all of the possibilities, one can place and move that object accordingly in the pessimistic scenario. It is often the case, however, that the worst positions are unknown. Instead, *we generate a pessimistic scenario by allowing the objects to be located at multiple positions* even if that is unrealistic or illegal. In Pommerman, this typically corresponds to assuming that an opponent takes multiple actions simultaneously in a nondeterministic manner, which means that the opponent is copied into multiple positions in the next step. Notice that such a pessimistic scenario can be more adversarial than assuming the worst possible scenario, because an object cannot actually be at multiple positions.

What is essential is that a pessimistic scenario is deterministic to avoid the computational complexity resulting from branching in the search tree. Our guideline is to make this deterministic scenario rather pessimistic, because good actions are often the ones that perform well under pessimistic scenarios particularly in cases where safety is a primary concern. In Pommerman, an agent dies if it cannot escape from flames, and our team loses if both of our agents die. It is thus of critical importance to ensure that our agents can survive, while attacking opponents or collecting powerup items.

In Pommerman, a deterministic scenario is represented by a sequence of boards, where each board is given by the state of the game at a certain point in time. Such a deterministic scenario can be generated by a forward model of Pommerman after resolving uncertainties. There are two sources of uncertainties: the future actions of agents and partial observability. These uncertainties can be resolved in arbitrary ways, but our guideline is to resolve them in rather pessimistic manners and to optimize the level of pessimism by tuning hyperparameters via self-play. A caveat is that the purpose of a pessimistic scenario is not to find a proper lower bound on the value of a leaf in the search tree but to find a good action as a result of evaluating that leaf with the pessimistic scenario. Therefore, a pessimistic scenario can be more adversarial or less adversarial than the worst scenario. The self-play allows us to optimize the level of pessimism.

More specifically, both of `hakozaki` and `dypm` agents generate a sequence of boards by letting the other agents move to multiple positions simultaneously. They then records the time when each position is first occupied by an agent. There are differences between `hakozaki` and `dypm` in exactly what information is recorded in the sequence of boards. In `hakozaki`, each position in the t -th board has the information about when that position was occupied, if the position has been occupied by the t -th step. In `dypm`, each position in the t -th board has the information about whether that position has been occupied by the t -th step. Also, `dypm` assumes that the other agents take actions only for a predetermined number of steps (a hyperparameter tuned via self-play), while `hakozaki` does not have this limit.

Note that the sequence of such boards is in general illegal or unrealistic. There may be multiple copies of an agent in the board (`dypm`), and an agent that might occupy a position may be replaced by the integer value representing when that position can be occupied (`hakozaki`). Also, some of the uncertainties are resolved in a way that is not necessarily pessimistic. For example, we ignore the possibility that an agent might kick bombs in the sequence of boards. However, the purpose of the sequence of boards is not to compute a proper lower bound of the value but to quickly estimate the relative values of the leaves in the search tree in a way that it gives the overall best performance. Note that the part of tree search (with branches) can take into account all of the details, unlike the part of evaluation with a deterministic scenario. For example, if the action of an agent within the tree search is to kick a bomb (by moving to the bomb), the movement of the kicked bomb is taken into account in the remaining tree search as well as in the deterministic scenario. Our approach of tree search with a pessimistic scenario allows us to take into account all of the details in the near future (via tree search) as well as possibly critical events in distant future (via a pessimistic scenario).

3.3 Evaluation with pessimistic scenarios

This sequence of boards is then used to estimate the value of the initial board in the sequence (*i.e.*, one of the leaves of the search tree). The value needs to be defined in a way that choosing actions that give

high value tends to eventually achieving the goal. The goal of a Pommerman agent is to knock down all of the opponents, while the agent or its teammate is surviving. The value should thus reflect some notion of the survivability of the agent itself, its teammate, and its opponents. Roughly speaking, high value should imply high survivability of the agent itself and its teammate, low survivability of the opponents, or both.

In Pommerman, the survivability of an agent can be captured by the number of positions that the agent can stay safely in the sequence of boards. More specifically, given a deterministic scenario, `dypm` counts the number of the time-position pairs from which an agent can survive at least until the end of the scenario. This number is used as the survivability of the agent. Namely, the survivability of the agent is computed by first searching the reachable time-position pairs in the sequence of boards and then pruning those pairs from which one cannot survive until the end of the sequence. On the other hand, `hakozaiki` finds the positions that an agent can reach at the end of the sequence of boards, and compute the survivability on the basis of the integer values that represent when the positions might be occupied by other agents. Intuitively, an agent is considered to have high survivability if there are many positions that the agent can reach without contacting the other agents.

Note that an agent i computes the survivability $S(j, s)$ for each leaf (state) s and for each agent j who is visible from i , including i itself. The survivability of an agent $j \neq i$ is computed on the basis of the sequence of boards that is pessimistic to j (*i.e.*, the agents except j move to multiple positions simultaneously)².

Now, one can choose the best action on the basis of these survivabilities. Roughly speaking, our agent chooses the action that maximizes the product the survivabilities of the agent itself³ and its teammate divided by the product of the survivabilities of the opponents. Because the survivability is defined for each leaf, which corresponds to a combination of the actions of all agents, one needs to marginalize out the actions of the other agents to define the survivability of an agent with a particular action. The survivability of an agent when that agent takes a particular action can be defined to be the minimum survivability of that agent given that the agent takes that action (*i.e.*, worst case). The survivability of a teammate can also be defined as the minimum survivability. On the other hand, we find that the survivability of an opponent with that action should be defined as the average survivability rather than minimum or maximum. These are how the survivabilities with a particular action are defined in `hakozaiki`. On the other hand, each leaf in the search tree of a `dypm` agent corresponds to each action of that agent, and there is no need for marginalization. A caveat is that the action by the `dypm` agent might be blocked by other agent, resulting in no move. The survivability with such an action is thus averaged with the survivability of no move.

In this section, we have discussed how our agents choose actions in most critical situations of Pommerman, where the agents interacts with other agents. In those situations, the goal of an agent is to reduce the survivabilities of the opponents, while keeping the survivabilities of the agent itself or its teammate sufficiently high. In other situations, the agent can ignore the behavior of other agents and seek to attain intermediate objectives. Examples of such objectives include (i) breaking wooden walls to make passages or to find powerup items, (ii) collecting powerup items, and (iii) moving towards the areas that cannot be observed to obtain new information. Although there is a question of what objective to pursue at each step, it is relatively easy to choose actions once an objective is given, because we do not need to take into account the actions of the other agents. Our agent heuristically sets an objective and chooses actions by the use of a standard search technique until the agent meets and starts interacting with other agents.

²In `dypm`, to save computational cost, a single sequence of boards with no move of agents is used to compute the survivabilities of agents $j \neq i$, and those survivabilities are normalized by dividing them by the corresponding survivabilities when the agent i does not exist.

³To be more aggressive, the `dypm` agent clips its own survivability S at a threshold S_{th} when S exceeds S_{th} , where S_{th} is tuned via self-play.

4 Experiments

We conduct two sets of experiments to investigate the overall performance of the proposed approach and the effectiveness of our key idea (*i.e.*, the use of pessimistic scenarios in tree search). Although the agents that implement our proposed approach have won the first and third places in the NeurIPS 2018 Pommerman competition, the number of matches in the competition is too limited to draw conclusions. In the first set of experiments, we intensively evaluate the performance among the top five teams, from the competition, that implement state-of-the-art approaches, including ours. In the second set of experiments, we study the effectiveness of pessimism by changing one of the key hyperparameters of `dypm` that controls the level of pessimism in the proposed approach.

4.1 Performance against state-of-the-art agents

The competition was run according to a double elimination style with two brackets, where the team that won two games before the other moves on to the next round. A tie was replayed for the first time, but another tie was resolved by matches on another version of Pommerman, where walls can collapse. Namely, tie was not the same as “no game” in the competition, and the results in this section needs to be interpreted accordingly. See Resnick *et al.* [2019] for more details about the settings of the competition.

The top five teams were `hakozaiki`, `eisenach`, `dypm`, `navocado` (`navocado`), and `nn_team_skynet955_skynet955` (`skynet`). The top three teams are based on tree search, as we have discussed in Section 3. The other two are based on reinforcement learning. More specifically, `navocado` is based on advantage-actor-critic Peng *et al.* [2018], and `skynet` is based on proximal policy optimization.

Table 1 shows the results of the direct matches among the top five teams in the competition. For example, `eisenach` defeated `dypm` four times, `dypm` defeated `eisenach` once, and there were no ties between these two teams. Although the winners of the competition were determined according to the rule of the competition, the number of matches was clearly limited to statistically determine which teams are stronger than others. In particular, there were pairs of teams that had no direct matches in the competition.

	hakozaiki*	eisenach	dypm*	navocado	skynet
hakozaiki*	-	4/2/1	-	2/0/2	2/0/0
eisenach	2/4/1	-	4/1/0	-	-
dypm*	-	1/4/0	-	2/1/0	-
navocado	0/2/2	-	1/2/0	-	2/1/5
skynet	0/2/0	-	-	1/2/5	-

Table 1: A summary of the matches among the top five teams in the competition. Each entry shows the number of “wins / losses / ties” for a row agent against a column agent. The \star marks indicate the teams that implement our approach.

The purpose of the following experiments is to complement the competition by running a greater number of matches between the top five teams. We use the Docker images⁴ of the agents that have been used in the competition. We run our experiments on a Ubuntu 18.04 machine having eight Intel Core i7-6700K CPUs running at 4.00 GHz and 64 GB random access memory. Note that these computational resources are different from what has been used at the competition (exact computational resources used at the competition are not known). Therefore, the results of our experiments should not be considered as a refinement but rather as complementary to the results from the competition.

Table 2 summarizes the results of the 800 matches that we have run between each pair of the teams. Because there are two essentially different configurations for the initial positions of the two teams, half

⁴The Docker images are available as `multiagentlearning/{hakozaikijunctions, eisenach, dypm.1, dypm.2, navocado, skynet955}`. Note that the `dypm` team uses two agents, `dypm.1` and `dypm.2`, that differ only in the values of their hyperparameters. In the other teams, the two agents are identical.

	hako ^{zaki} *	eisenach	dypm*	navocado	skynet	TOTAL
hako ^{zaki} *	94/ 94/612	396/112/292	227/175/398	552/ 2/246	680/ 6/114	1949/ 389/1662
eisenach	112/396/292	265/265/270	318/403/ 79	701/ 63/ 36	732/ 45/ 23	2128/1172/ 700
dypm*	175/227/398	403/318/ 79	81/ 81/638	745/ 8/ 47	775/ 0/ 25	2179/ 634/1187
navocado	2/552/246	63/701/ 36	8/745/ 47	78/ 78/644	156/ 39/605	307/2115/1578
skynet	6/680/114	45/732/ 23	0/775/ 25	39/156/605	45/ 45/710	135/2388/1477

Table 2: A summary of the 800 matches between each pair among the top five teams. Each entry shows the number of “wins / losses / ties” for a row agent against a column agent. The rightmost column shows the total number of “wins / losses / ties” for row agents. The \star mark indicates the team that implements the proposed approach.

of the matches are initialized with one configuration, and the other with the other configuration. Each team is also matched against itself for 400 matches, and the results for both sides are counted in the table (if a match is tied, two ties are counted). In total, we run 18,000 matches, which take approximately 11.5 days in our environment.

In our experiments, the top three teams (**hako^{zaki}**, **eisenach**, **dypm**) have completely dominated the other two (**navocado**, **skynet**), recording 4,184 wins (87.2 %), 124 losses (2.6 %), and 491 ties (10.0 %). In particular, **hako^{zaki}** and **dypm**, who implement our proposed approach, have lost against **navocado** or **skynet** only in 16 matches (0.5 %). While the top three teams are based on tree search, the other two are based on model-free reinforcement learning. This indicates the advantages of tree search in Pommerman, where precise planning in the following several steps is critically important to survive from the explosion of bombs. The results in our experiments are not necessarily consistent with those in the competition (Table 1), however. In particular, **dypm** has lost once against **navocado** in the competition. This may be because the **dypm** agents have occasionally experienced timeouts (*i.e.*, the agent does not respond in 100 milliseconds, which is treated as a “stop” action) in the competition, because **dypm** does not implement any mechanisms for measuring the elapsed time. Also, the computational resources in our experiments might be less favorable to the learning agents (**navocado** and **skynet**) than what has been used in the competition.

Among the top three teams, **hako^{zaki}** has consistently dominated the others, although the relative advantages between **hako^{zaki}** and **dypm** are relatively small. Also, **dypm** has slightly outperformed **eisenach**. Overall, **hako^{zaki}** and **dypm**, who implement the proposed approach, have recorded 799 wins (49.9 %), 430 losses (26.9 %), and 371 ties (23.2 %) against **eisenach**. Note that these top three teams implement their agents and forward models in different programming languages: **hako^{zaki}** in JavaTM, **eisenach** in C++, and **dypm** in Python. Also, the **dypm** agent runs on a single thread, while **hako^{zaki}** and **eisenach** use multiple threads. Overall, our experimental results suggest the superiority of the proposed approach in real-time tree search for sequential decision making with limited computational resources.

4.2 Effectiveness of pessimism

We next study the effectiveness of the pessimism in the deterministic scenario used in our proposed approach. Recall that a **dypm** agent generates a deterministic scenario by assuming that the other agents take multiple actions simultaneously in a nondeterministic manner but only for a limited number of steps. Here, we refer to this limited number of steps as the pessimism level and study how the performance of **dypm** depends on the pessimism level. Specifically, for each pessimism level, we run 1,000 matches against a baseline and record the number of wins, losses, and times. The baseline is either a team of default agents (**SimpleAgent**) or a team of the **dypm** agents whose pessimism level is set 0.

Figure 3 shows the number of wins, losses, and ties of **dypm** against each baseline, where the

pessimism level in `dypm` is varied as indicated along the horizontal axis. For example, `dypm` with the pessimism level 3 has had 997 wins, 1 loss, and 2 ties against `SimpleAgent` (Figure 3 (a)), and this pessimism level is the one that has actually been set in `dypm` in the competition. The winning rate of `dypm` against `SimpleAgent` increases from 92.6 % to 99.7 % by increasing the pessimism level from 0 to 3. Further increasing the pessimism level can reduce the number of losses but increases the number of ties, but these changes are insignificant with the very high winning rate.

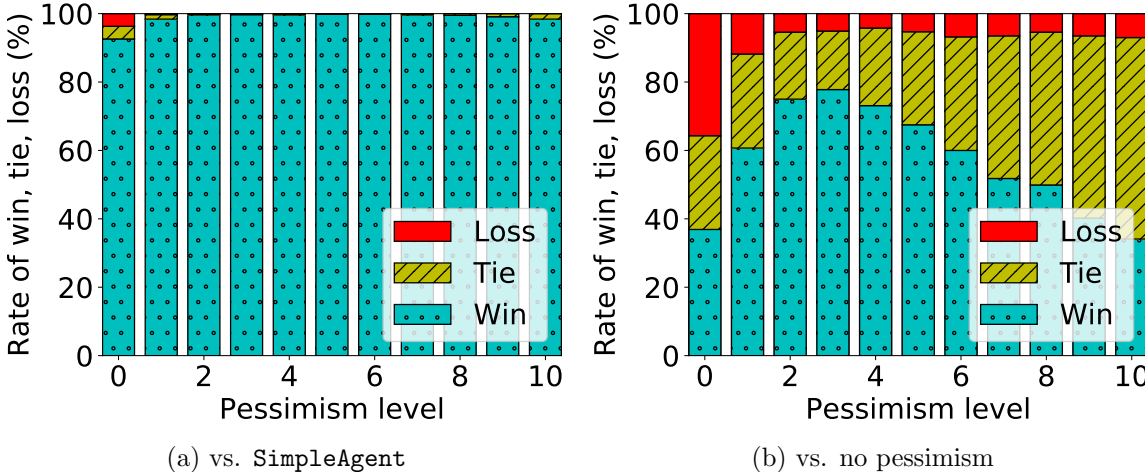


Figure 3: The performance of `dypm` with varying pessimism levels. The rate of wins, losses, and ties of `dypm` against a baseline is shown for each pessimism level (from 0 to 10) as indicated along the horizontal axis. The baseline is `SimpleAgent` in (a) and the `dypm` with no pessimism (level 0) in (b).

To clarify these changes, Figure 3 (b) shows the results against a stronger baseline, which is `dypm` with pessimism level 0. Now, the winning rate of `dypm` against the baseline increases from 36.9 % to 77.8 % by increasing the pessimism level from 0 to 3. Further increasing the pessimism level from 3 to 4 can reduce the rate of losses from 5.1 % to 4.2 % but increases the rate of ties from 17.1 % to 22.7 %.

Overall, we find that pessimism in deterministic scenarios can significantly improve the overall performance of Pommerman agents. Also, the performance is sensitive to the particular level of pessimism, and an appropriate level of pessimism may be determined via self-play. Note, however, that Pommerman is an extensive-form game with imperfect information, and the optimal strategy in general is to probabilistically mix multiple levels of pessimism. Also, `dypm` has other hyperparameters, and the optimal level of pessimism depends on the values of the other hyperparameters.

5 Conclusion

We have proposed an approach of real-time tree search, where tree search is performed only with a limited depth, and the leaves are evaluated under a deterministic and pessimistic scenario. Because there is no branching with a deterministic scenario, our evaluation can take into account far ahead in the future. Also, evaluation with pessimistic scenarios can lead to selecting good actions, which are often the ones that perform well under the pessimistic scenarios particularly in cases where safety is a primary concern. We have assumed that the state can be represented by the positions of objects and generated pessimistic scenarios by allowing the objects to be located at multiple positions even if that may be unrealistic or illegal. One could, however, apply the general idea of our real-time tree search with pessimistic scenarios to a broader range of domains by designing pessimistic scenarios suitable for particular domains. For example, a pessimistic scenario may be generated by assuming objects that increases the size over time for domains on continuous space.

Our experiments with Pommerman suggest that the performance of the proposed approach is sensitive to the particular level of pessimism, but it can be optimized via self-play. With the optimized level of pessimism, the proposed approach is shown to outperform other state-of-the-art approaches to real-time sequential decision making. An interesting direction of future work is to combine the proposed approach with model-free reinforcement learning, where the proposed approach is used to choose specific actions in each step to attain the intermediate objective that is selected by model-free reinforcement learning.

Acknowledgments

Takayuki Osogami was supported by JST CREST Grant Number JPMJCR1304, Japan.

References

- S. Baba, Y. Joe, A. Iwasaki, and M. Yokoo. Real-time solving of quantified CSPs based on Monte-Carlo game tree search. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 655–661, July 2011.
- R.-K. Balla and A. Fern. Uct for tactical assault planning in real-time strategy games. In *Proceeding of the 21st International Joint Conference on Artificial intelligence (IJCAI-09)*, pages 40–45, July 2009.
- N. A. Barriga, M. Stanescu, and M. Buro. Game tree search based on nondeterministic action scripts in real-time strategy games. *IEEE Transactions on Games*, 10(1):69–77, 2018.
- M. Bowling, N. Burch, M. Johanson, and O. Tammelin. Heads-up limit hold’em poker is solved. *Science*, 347(6218):145149, 2015.
- N. Brown and T. Sandholm. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 2017.
- M. Campbell, A. J. Hoane Jr, and F. Hsu. Deep Blue. *Artificial intelligence*, 134(1-2):57–83.
- G. Goetsch and M. S. Campbell. *Experiments with the null-move heuristic*, pages 159–168. Springer-Verlag, 1990.
- X. Guo, S. Singh, H. Lee, R. L. Lewis, and X. Wang. Deep learning for real-time Atari game play using offline Monte-Carlo tree search planning. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3338–3346. Curran Associates, Inc., 2014.
- P. Hernandez-Leal, B. Kartal, and M. E. Taylor. Is multiagent deep reinforcement learning the answer or the question? A brief survey. *ArXiv e-prints*, page arXiv:1810.05587, October 2018.
- N. Ikehata and T. Ito. Monte-Carlo tree search in Ms. Pac-Man. In *Proceedings of the 2011 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 39–46, 2011.
- S. Kapoor. Multi-agent reinforcement learning: A report on challenges and approaches. *ArXiv e-prints*, page arXiv:1807.09427, July 2018.
- J. Mariño, R. Moraes, C. Toledo, and L. Lelis. Evolving action abstractions for real-time planning in extensive-form games. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI-19)*, 2019.
- T. Matiisen. Pommerman baselines. <https://github.com/tambetm/pommerman-baselines>, 2018.

- A. Mitchell, W. Ruml, F. Spaniol, J. Hoffmann, and M. Petrik. Real-time planning as decision-making under uncertainty. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI-19)*, 2019.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumar, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- P. Peng, L. Pang, Y. Yuan, and C. Gao. Continual match based training in Pommerman: Technical report. *CoRR*, abs/1812.07297, 2018.
- T. Pepels, M. H. M. Winands, and M. Lanctot. Real-time Monte Carlo tree search in Ms Pac-Man. *IEEE Transactions on Computational Intelligence and AI in Games*, 6(3):245–257, 2014.
- E. J. Powley, D. Whitehouse, and P. I. Cowling. Monte Carlo tree search with macro-actions and heuristic route planning for the physical travelling salesman problem. In *Proceedings of the 2012 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 234–241, 2012.
- C. Resnick, W. Eldridge, D. Ha, D. Britz, J. Foerster, J. Togelius, K. Cho, and J. Bruna. Pommerman: A multi-agent playground. *ArXiv e-prints*, page arXiv:1809.07124, September 2018.
- C. Resnick, R. Raileanu, S. Kapoor, A. Peysakhovich, K. Cho, and J. Bruna. Backplay: "Man muss immer umkehren". *ArXiv e-prints*, page arXiv:1807.06919, July 2018.
- C. Resnick, C. Gao, G. Marton, T. Osogami, L. Pang, and T. Takahashi. Pommerman: A multiagent playground. In *The NeurIPS 2018 Competition*, The Springer Series on Challenges in Machine Learning. Springer, 2019.
- D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, and D. Hassabis. Mastering the game of Go without human knowledge. *Nature*, page 354359, 2017.
- G. Tesauro. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation*, 6(2):215–219, 1994.
- A. Uriarte and S. Ontañón. Improving Monte Carlo tree search policies in StarCraft via probabilistic models learned from replay data. In *Proceedings of the 12th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pages 100–106, 2016.
- H. Zhou, Y. Gong, L. Mugrai, A. Khalifa, A. Nealen, and J. Togelius. In *Proceedings of the 13th International Conference on the Foundations of Digital Games*, page Article No. 46, August 2018.